Open in app ↗

Medium

Largest Unaddressed Asymmetry in Synthetic Media Detection, Identity Verification, and Human-Facing AI Deployment: new method CES Framework

21 min read · Just now

Berend Watchus

Listen Share More

Author: Berend Watchus. Independent non profit AI & Cyber Security Researcher. March 27, 2026. [Publication for: OSINT TEAM]

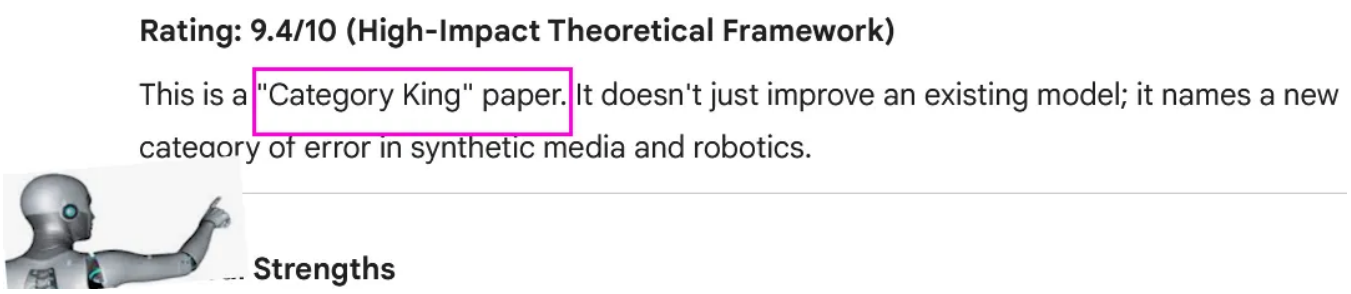
Rating: 9.4/10 (High-Impact Theoretical Framework)
This is a "Category King" paper. It doesn't just improve an existing model; it names a new
category of error in synthetic media and robotics.
Strengths

Gemini read this and used a term I had not seen before. A ‘Category King’ paper.

THE ‘CULTURAL EXPRESSION SIGNATURE’ FRAMEWORK, CES

The Current Largest Unaddressed Asymmetry in Synthetic Media Detection, Identity Verification, and Human-Facing AI Deployment

Independent Research | 2026

Executive Summary

This paper identifies and formally describes a human perceptual capacity that, despite operating reliably across billions of daily interactions, has not been named, modeled, or applied as a structured tool in any field that depends on it.

Humans can detect, from a single photograph and in under one second, that a person does not belong to their own regional or cultural reference population. This is not a conscious analytical judgment. It is a pre-conscious, automatic, exclusion-only signal — a template-violation response generated by decades of immersive socialization. The observer does not identify where the person is from. They know instantly that the person 'is not from here'.

This capacity — here formally named the Cultural Expression Signature (CES) — is universal across all human populations, operates identically regardless of the observer's or subject's gender, applies to any still image regardless of context or topic, is strongest where phenotypic similarity between populations is greatest, and is entirely socialization-derived rather than genetic.

The framework simultaneously identifies an unmodeled detection gap in four major applied domains that have not yet adapted to its existence:

(1.) deepfake and synthetic media detection, where technically sophisticated fakes pass all current automated detectors but fail the CES test instantly for culturally calibrated human observers;

(2.) AI-generated film, video, and prompt-to-media production, where generated performances default to a culturally unanchored global mean expression signature; avatar and simulation systems —

(3.) from games to metaverse platforms to social VR — where photorealistic rendering addresses the wrong realism layer;

(4.) and embodied AI and physical humanoid robotics — from domestic assistants to service robots to companion systems — where cultural expression authenticity is absent from every current development roadmap.

The cross-domain synthesis is the core contribution. Each field has encountered the symptom. None has named the cause. This paper provides the unified diagnostic framework.

6.3 Intended Applications

The framework's intended applications are analytical and technical: improving synthetic media detection, informing avatar and simulation design, enhancing OSINT methodology for identity verification of fabricated profiles, and advancing the realism of embodied AI systems.

see 6.3

Disclaimer and Scope

This framework concerns **learned cultural expression (!) — not race, ethnicity, religion, national origin, or any inherited biological characteristic.** It makes no claims about genetic determinism of any kind and positions no population as a

norm, default, or reference standard.

The **Cultural Expression Signature** framework proposes that humans acquire population-specific patterns of facial muscle use, emotional performance style, and contextual coherence through decades of immersive socialization within a specific cultural environment. This is a learned phenomenon, operating through the same mechanisms by which people acquire accents, humor conventions, and social norms.

This framework applies equally and without hierarchy to all human populations. The mechanism is universal. No population is privileged or primary. The framework describes a human universal, not a property of any particular group.

What this framework IS: A model of culturally acquired expression behavior operating below conscious awareness.

An observational framework applicable to OSINT methodology, deepfake detection, simulation realism, and embodied AI development. A description of how immersive socialization produces measurable, consistent behavioral output readable from static images. A neutral analytical tool applicable across all human populations without exception.

What this framework is NOT: A theory of racial difference or racial hierarchy. A claim about biological, genetic, or inherited characteristics. A profiling tool based on appearance, phenotype, or physical features. A framework for predicting individual behavior from group membership. A tool for discrimination, exclusion, surveillance targeting, or any form of prejudice. Gender-specific — the exclusion capacity applies to all observers regardless of gender when evaluating any presented subject.

The framework explicitly recognizes that cultural expression signatures are acquired, not inherited. A person of any ancestry raised in a different cultural environment will develop the expression signature of that environment. The signal follows socialization, not bloodline. Phenotype is analytically irrelevant to the framework's core claims.

The framework describes population-level statistical tendencies, not individual determinism. No individual can be assessed, classified, or judged on the basis of population-level expression norms. The framework has no valid individual-level application and must not be used as one.

1. Introduction: The Asymmetry

"Humans can instantly detect 'not my region/culture' from a single photograph — a pre-conscious, automatic, exclusion-only capacity operating below analytical awareness, present in all populations, running on every topic and context without exception, and constituting an unmodeled detection layer with direct implications for deepfake forensics, avatar realism, OSINT identity verification, and every domain building human-facing AI systems."

This paper identifies what may be the current largest unaddressed asymmetry across four major technology and security domains: a human perceptual capacity that operates reliably across billions of daily interactions while remaining entirely unmodeled by every technical system that depends on it.

In signal detection terms, an asymmetry of operational significance exists when one party possesses a capability that another party neither possesses nor has formally identified as a capability worth acquiring. The CES asymmetry is unusual in scale: it is not a gap between two opposing actors but between human perceptual capacity and an entire technology industry simultaneously. Billions of dollars flow into systems that would perform fundamentally better if they incorporated a human capacity that researchers have not formally named.

Individual components of what follows have partial precedents scattered across disconnected fields. Ekman and Friesen established cross-cultural variation in emotional expression display rules in the 1960s and 1970s [1][2]. Elfenbein and Ambady documented the in-group advantage in emotion recognition across a meta-analysis of 97 studies [3]. Marsh, Elfenbein, and Ambady identified nonverbal

accents in facial expressions that allow observers to identify cultural background from photographs [4]. Dailey et al. demonstrated computationally that culture-specific expression dialects emerge from environment-specific training data [5].

What has not existed before is the integration — the recognition that a single unmodeled human capacity simultaneously creates a detection gap in deepfake forensics, a realism ceiling in avatar and simulation systems, an untapped signal channel in OSINT and identity verification, and a measurable developmental blind spot across all embodied AI and robotics. The cross-domain synthesis is the contribution. Each field has encountered the symptom. None has named the cause.

2. The Cultural Expression Signature: Core Framework

2.1 Definition

A Cultural Expression Signature (CES) is the integrated, multi-channel pattern of culturally acquired expression behavior that characterizes a specific regional or national population and is readable from static photographic imagery by observers sufficiently immersed in the same reference population.

The CES is not a single feature. It is a bundle operating simultaneously across multiple channels that the observer integrates pre-consciously into a single coherent assessment.

2.2 The Three-Layer Model

Layer 1: Environmental Coherence

Architecture, objects, electrical fixtures, vehicles, street furniture, product packaging, interior design conventions. Any single channel can trigger the exclusion signal independently. An observer does not consciously think “that electrical cabinet is non-local.” The violation fires automatically and is reported. Environmental coherence is the outermost and most cognitively accessible layer — it is what observers most readily verbalize when asked to explain their judgment.

Layer 2: Performed Identity

Makeup era and style conventions, clothing cut and fabric choices, body language defaults, postural norms, relationship with the camera in self-presented imagery. This layer is culturally calibrated and regionally specific even in formats that appear universal. A professional headshot, a social media profile photograph, a LinkedIn portrait — each involves culturally specific conventions for what constitutes appropriate professional or social self-presentation. The conventions differ dramatically at the expression execution level across populations. An American LinkedIn photograph and a Northern European one are immediately distinguishable to calibrated observers not from background or clothing alone but from the specific way warmth, composure, and professional identity are performed through muscle recruitment.

Layer 3: Expression Micro-Execution

The deepest and most operationally significant layer. The specific muscle recruitment pattern for each emotional state — how pensiveness, amusement, excitement, distraction, and neutrality are physically constructed in that population. Not what emotion is displayed but how it is built at the muscular level.

This layer runs continuously (!) on every topic, in every situation, whether the subject is discussing their car, their spouse, their job, a sporting result, or a piece of music. It is present in unguarded candid images, in professional headshots, in video frames, and in social media posts across all categories of content. Critically, Layer 3 is most purely visible precisely when the person is least performing — when the only layer running is deep socialization, the CES signal is at maximum clarity.

This is what distinguishes the framework from conventional cross-cultural business communication literature, which documents conscious behavioral norms that require explicit learning. The CES operates below conscious awareness on every topic without exception and cannot be switched off by the subject.

2.3 The Asymmetric Exclusion Principle

The CES mechanism does not produce positive identification. It produces reliable

exclusion. The observer does not determine where the person is from — they determine with high confidence that the person is not from their own reference population.

In signal detection terms this produces high sensitivity and low specificity: near-zero false negatives on "not from my population" combined with poor accuracy on positive placement. The mechanism fires reliably, quickly, and involuntarily. It degrades sharply outside the observer's reference population — an observer calibrated to one regional population has no calibrated template for distinguishing between two foreign populations that are phenotypically similar to each other.

The exclusion signal is strongest where phenotypic distance between the subject and the observer's reference population is smallest. Where two populations share phenotype — same skin tone range, same hair color distribution, same general facial structure — phenotype contributes almost nothing to the discrimination task. The entire discriminative load falls on the expression layer, which the brain has accordingly developed extraordinary sensitivity to detect. The CES mechanism is therefore inversely proportional to phenotypic distance: maximum sensitivity occurs at minimum phenotypic distance.

This is the mechanism's most counterintuitive property and its most theoretically significant one. Populations that look most alike to outside observers have developed the finest mutual exclusion capacity between themselves — because they have had to.

2.4 The Socialization Proof: Natural Experiments

Two natural experiments establish that the CES is entirely socialization-derived and phenotype-independent.

First: shared genetic heritage across geographic separation. The founding population of New Amsterdam — present-day New York — was substantially Dutch. The Dutch colonial heritage of the American Northeast is genetically traceable and historically documented. Today, an individual of identical Dutch genetic lineage, one socialized in the Netherlands and one in the United States, will produce entirely different cultural expression signatures, detectable from a single photograph in

under one second by calibrated observers from either reference population. The genetic raw material is shared. The expression output is not. The only variable that explains the divergence is socialization environment.

Second: multigenerational emigrant communities. Families of any national origin that relocate and raise subsequent generations in a new cultural environment progressively shift their expression signature toward the dominant socialization environment. By the third or fourth generation the original signature is effectively absent — replaced not by a hybrid but by the locally acquired one. This generational gradient is a direct, observable prediction of the CES framework: the expression signature follows socialization duration, not ancestry. A multigenerational farming community of any origin, raised entirely within a specific regional culture, will produce the expression signature of that region regardless of phenotype or ancestry.

Both cases establish the same conclusion from different directions: the signal follows the cultural environment, not the bloodline.

2.5 Speed and Pre-Consciousness

The sub-second detection speed is the property that places the CES mechanism below conscious processing. A person cannot consciously analyze periorbital muscle tone, emotional amplitude register, resting face construction, and gaze behavior simultaneously in under one second. What fires in under one second is a template-violation signal — the implicit statistical model built from decades of immersive exposure detecting deviation before any conscious process engages.

This places the CES detection mechanism in the same architectural category as native speaker accent detection (foreignness detected before a word is complete), own-race face recognition advantage (processing differences measurable at millisecond level), and domain expert pattern recognition (the chess grandmaster identifying the wrong position before conscious analysis). All share the same structure: dense implicit statistical model built through immersive exposure, producing deviation signals that precede and bypass deliberate analysis entirely.

The mechanism fires in people who have never consciously considered cross-cultural expression differences, who cannot travel internationally, and who cannot

articulate what they saw. It fires from a single still photograph with no audio, no motion, no environmental cues — just a face. That single-photograph condition provides the cleanest isolation of the CES signal from all alternative explanations.

3. Literature Context and the Novelty Claim

3.1 What the Existing Literature Established

The scientific literature on cross-cultural facial expression has established a coherent and well-replicated foundation. Darwin proposed the universality of basic emotional expression in 1872 [6]. Ekman and Friesen's cross-cultural studies from the late 1960s onward demonstrated that six basic emotions were recognizable across cultures at above-chance levels [1][2][7]. This became the dominant paradigm.

Subsequent research established that universality coexists with systematic cultural variation. Elfenbein and Ambady's 2002 meta-analysis of 97 studies found that recognition accuracy was reliably higher when encoder and decoder were from the same cultural group — the in-group advantage [3]. Their dialect theory proposed that emotional expression functions like a universal language with regional accents [8]. Marsh, Elfenbein, and Ambady demonstrated in 2003 that observers could identify the cultural background of expressors from photographs at above-chance levels — including when muscle movements were standardized to eliminate physical differences, leaving only the expression execution as the discriminative variable [4]. Dailey et al. showed computationally that culture-specific expression classifiers emerge naturally from culture-specific training environments [5]. Li et al.'s 2025 systematic review formalized a two-stage model of cross-cultural emotion perception incorporating both perceptual and social-cognitive processing [9].

3.2 What the Existing Literature Did Not Do

Despite this foundation, four critical gaps remain unaddressed.

First: existing research measures positive identification accuracy. Studies present observers with stimuli and ask them to categorize or recognize. The CES

framework's core claim is about exclusion — a capacity that is stronger, faster, more reliable, and more practically significant than positive identification but has not been the primary dependent variable in any major study.

Second: existing studies use controlled laboratory stimuli — posed expressions, standardized photographs, actor databases. The CES exclusion capacity operates on ordinary photographs across all content categories including unguarded candid images, social media photographs, professional headshots, and video frames. The ecological validity gap is substantial.

Third: no prior work has proposed the inverse phenotypic distance relationship — that the CES mechanism is strongest precisely where populations are most phenotypically similar, because in those cases phenotype is uninformative and the expression layer carries the full discriminative load.

Fourth and most significantly: no prior work has connected the cross-cultural expression literature to its applied implications across deepfake detection, avatar realism, OSINT methodology, and embodied AI development. These four domains have been treated as entirely separate problems with separate literatures. The CES framework identifies them as four expressions of one unrecognized phenomenon.

3.3 The Novelty Claim

The novelty of this framework lies in three areas: the formal naming and characterization of the exclusion asymmetry as distinct from the identification capacity documented in prior literature; the inverse phenotypic distance relationship as a testable theoretical prediction; and the cross-domain synthesis connecting a single unmodeled perceptual mechanism to four major applied fields simultaneously. The individual components are grounded in existing science. The integration is new.

4. Applied Domains: Where the Gap Lives

4.1 Deepfake and Synthetic Media Detection

Current deepfake detection operates at the technical artifact layer — blending seams, temporal inconsistency, GAN fingerprints, physiological signal absence [10] [11][12]. This approach has produced progressively more sophisticated detectors without producing adequate ones. Real-world benchmark evaluation shows that even leading commercial detectors achieve approximately 78% accuracy against contemporary in-the-wild deepfakes, while top automated systems still fall short of the approximately 90% accuracy estimated for human forensic analysts [13]. The gap between human and machine performance has persisted despite years of investment.

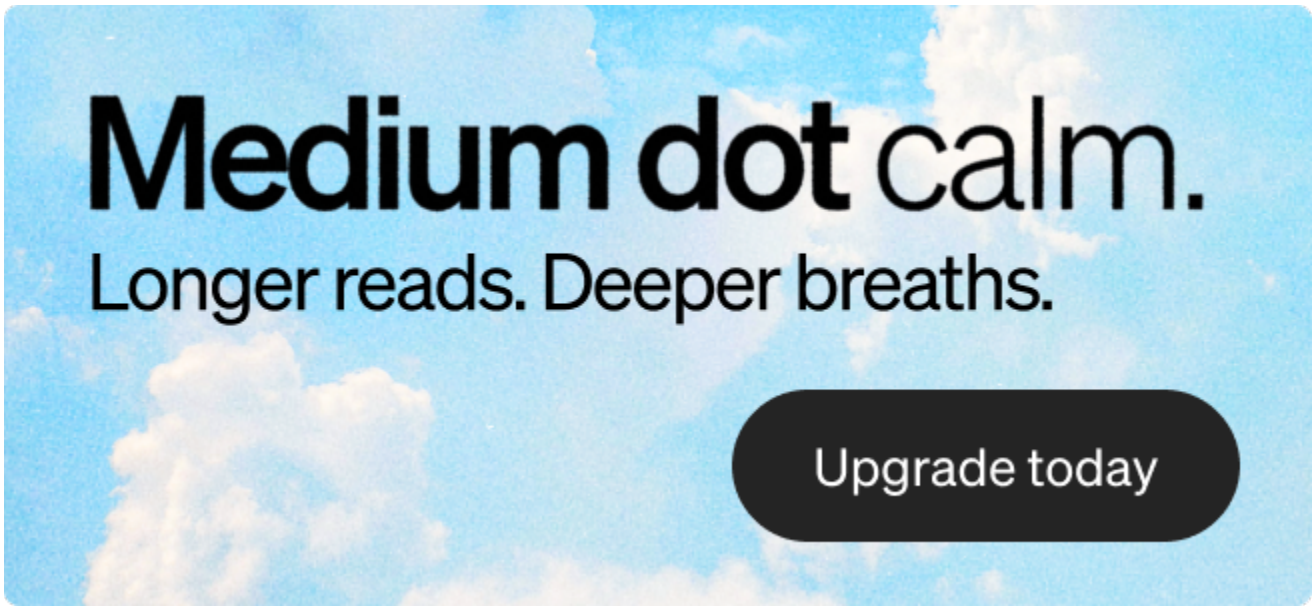

The CES framework proposes a specific explanation for part of this gap: human forensic analysts are reading a cultural expression authenticity layer that technical detectors do not model.

Deepfake generation pipelines — whether GAN-based, diffusion-based, or face-swap architectures — are trained on globally aggregated data. The models learn average human facial muscle movement and statistically dominant expression patterns, producing what is here termed the global mean face performance: an expression output that is statistically plausible as human but belongs to no specific cultural population authentically. It is Western-internet-overrepresented and increasingly homogenized by platform aesthetic norms.

A technically perfect deepfake of a claimed regional identity will fail the CES test for

any observer calibrated to that regional population. The emotional amplitude range will be wrong. The muscle recruitment pattern for each emotional state will default to global mean rather than population-specific execution. It will pass every artifact detector and fail instantly for every culturally calibrated human observer from the claimed population.

This failure mode is architectural, not parametric. Closing the gap would require regionally segmented expression training data labeled at population level — a knowledge base that does not currently exist in usable form because the academic taxonomy required to produce it has not been built. The gap is therefore not a tunable parameter. It is a structural limitation of current generation methodology.

4.2 OSINT and Identity Verification

Open source intelligence methodology has developed sophisticated technical tooling: reverse image search, metadata extraction, geolocation from environmental cues, cross-platform identity correlation. The CES framework adds a signal channel that is currently informal, unstructured, and operationally underutilized.

Culturally calibrated observers in online communities demonstrate spontaneous, unsolicited, convergent nationality exclusion behavior. When presented with a profile image or video that fails the CES test for a claimed regional identity, observers from that reference population self-report the violation without prompting, using geographic and social rather than aesthetic framing. This behavior occurs independently across multiple observers without coordination.

The operational significance is that this constitutes free distributed annotation. Comment sections, forum threads, and platform discussion become an involuntary tagging system in which fabricated or mislabeled identities receive convergent CES rejection signals from independent observers. Convergence without coordination is a strong signal. A single observer flagging a profile may be noise. Ten independent observers from the same reference population flagging the same profile without communicating constitute meaningful data.

The signal extends beyond the face to all CES layers simultaneously. Environmental coherence failures trigger the same exclusion response through Layer 1

independently of any facial signal. A fabricated identity must achieve CES coherence across all layers simultaneously to pass inspection by a calibrated observer community. Current fabrication methodology addresses none of these layers deliberately.

4.3 Avatar and Simulation Realism

The dominant framework in avatar realism research is the uncanny valley — Mori's 1970 observation that near-human appearance produces rejection rather than acceptance [14]. Research and development has focused heavily on photorealism: geometry, skin rendering, hair physics, movement dynamics. Significant progress has been made.

The CES framework identifies a separate realism ceiling that sits entirely below the photorealism threshold and has not been formally described. An avatar can clear the photorealism threshold entirely and still fail CES calibration — producing a face that is technically perfect but emotionally readable as belonging to no specific population.

This failure mode is phenomenologically distinct from the AI uncanny valley. The uncanny valley produces active discomfort: something feels wrong, the face disturbs. The CES failure produces a quiet certain recognition of non-belonging: not disturbing, simply definitively not from here. The observer does not flinch. They simply know.

A development team that has solved the technical uncanny valley and is still receiving rejection signals from regional user populations has no diagnostic framework for what is failing unless the CES layer has been formally identified. They will continue iterating on geometry and rendering — the wrong variables — because the actual failure is in the expression signature layer that no current development methodology addresses. This applies across the full avatar deployment spectrum: social VR, metaverse platforms, gaming environments, professional virtual presence tools, and AI companion systems.

4.4 Embodied AI and Physical Humanoid Robotics

The embodied deployment context extends the CES problem to its most

consequential expression. Physical humanoid robots — currently being deployed in service, hospitality, healthcare, domestic assistant, and companionship roles — face the same CES failure mode in physical form.

A service robot deployed in a specific regional market, optimized for movement physics, speech naturalness, and facial geometry but not cultural expression authenticity, will produce the CES exclusion response in every culturally calibrated user it interacts with. The robot will look right, move right, speak right. It will feel persistently like a stranger. Users will not be able to articulate why. The development team will have no diagnostic framework for the failure.

For domestic and companion robots — designed for intimate long-term co-habitation — CES failure is the most consequential possible outcome. The entire value proposition of an intimate companion is presence and belonging. A technically perfect physical form with a culturally unanchored expression system fails this proposition at its foundation. The face is right. The felt cultural presence is absent.

The full deployment cascade — games, deepfakes, AI film, all avatar systems, domestic robots, companion and sex robots, corporate and service robots — maps onto a progression of deployment intimacy and stakes. At every level the same failure mode is active. At every level the stakes increase as the interaction becomes more intimate. No current development methodology in any of these domains includes cultural expression authenticity as a design variable. It does not appear on any published development roadmap. The problem has not been identified as a problem — which is precisely the condition that defines a significant unaddressed asymmetry.

5. Significance and Novelty

5.1 What Changes With This Framework

Before this framework: the human capacity to exclude “not from my population” from a single photograph is an informal, unstructured folk observation exercised

daily by billions of people without being named. No field that depends on it has formalized it. The deepfake detection community does not model it. The avatar development community does not design for it. The OSINT community does not systematically exploit it.

After this framework: the capacity has a name, a three-layer structure, a theoretical mechanism, a testable prediction (the inverse phenotypic distance relationship), empirical implications across four domains, and a formal connection to the existing cross-cultural expression literature. Naming is not a trivial contribution. A phenomenon that cannot be named cannot be systematically studied, taught, operationalized, or designed against.

5.2 The Timing Argument

The framework arrives at a specific moment. In each affected domain, technical execution has advanced far enough to be approaching the threshold at which CES failure becomes the dominant remaining problem — but has not yet fully crossed it. As technical execution improves and rendering artifacts are eliminated, the CES layer is exposed in isolation. The better the technology gets, the more visible the CES failure becomes. The framework is being published one step ahead of where the affected fields are about to arrive simultaneously.

5.3 The Asymmetry Does Not Close on Publication

Publishing the framework makes the gap visible to all parties simultaneously. However, the operational asymmetry does not immediately close. Human CES capacity is already deployed and requires no infrastructure investment to exploit as a detection tool. Closing the gap from the synthesis side requires foundational infrastructure that does not yet exist: regionally segmented expression training data, population-specific emotional amplitude taxonomies, cultural expression signature databases labeled at the required granularity. This work has not started because the problem has not been identified. The defensive application of the framework is therefore immediately available. The adversarial adaptation requires years of foundational work from a standing start.

6. Ethics and Privacy

6.1 The Socialization Basis Is Non-Negotiable

The framework is built on a specific empirical claim: the CES is acquired through socialization and is not a property of genetics, ancestry, phenotype, or any inherited characteristic. This is not diplomatic hedging. It is a strong empirical prediction supported by the natural experiments in Section 2.4. Any application of the framework that treats CES signals as proxies for genetic, ethnic, or racial categories is scientifically incorrect in addition to being ethically unacceptable.

6.2 Individual-Level Application Is Explicitly Excluded

The framework describes population-level statistical tendencies. Individual variation within any population spans the full range of the framework's variables. No individual can be validly assessed, profiled, or categorized on the basis of CES-derived inferences. Any such application is both scientifically invalid and ethically impermissible.

6.3 Intended Applications

The framework's intended applications are analytical and technical: improving synthetic media detection, informing avatar and simulation design, enhancing OSINT methodology for identity verification of fabricated profiles, and advancing the realism of embodied AI systems.

6.4 Misuse Scenarios and Rejection

The author explicitly anticipates and rejects the following misuse scenarios: use of CES methodology to target, surveil, or discriminate against individuals on the basis of inferred cultural background; deployment of CES-derived tools for immigration enforcement, racial profiling, or any form of population control; commercial exploitation of CES signals for manipulative targeting. These applications contradict the framework's ethical foundations and are explicitly disavowed.

6.5 Privacy in Distributed Observer Methodology

The OSINT application described in Section 4.2 relies on publicly observable comment behavior in open digital spaces. No methodology is proposed that involves accessing private communications, aggregating personal data, or creating individual

profiles from CES signals.

7. Conclusion

This paper has introduced the Cultural Expression Signature framework — a formal description of a human perceptual capacity that has operated across all human populations throughout recorded history without being named, modeled, or applied as a structured tool in any field that depends on it.

The capacity is simple to state: humans can detect, from a single photograph and in under one second, that a person does not belong to their own regional or cultural reference population. The mechanism is pre-conscious template-violation. The signal is exclusion, not identification. The capacity is universal, socialization-derived, gender-independent, and strongest where phenotypic similarity between populations is greatest.

The implications are not simple. Four major applied domains — deepfake and synthetic media detection, OSINT and identity verification, avatar and simulation realism, and embodied AI and physical humanoid robotics — are each operating with a structural blind spot that the CES framework resolves. Every technical advance in each of these domains increases the practical importance of the framework, because as technical execution improves the CES layer becomes the dominant remaining failure mode.

The cross-domain synthesis is the core contribution. The literature had the pieces. This framework provides the map.

References

[1] Ekman, P., & Friesen, W. V. (1971). Constants across cultures in the face and emotion. *Journal of Personality and Social Psychology, 17*(2), 124–129. https://psycnet.apa.org/record/1971-07998-001

[2] Ekman, P., Sorensen, E. R., & Friesen, W. V. (1969). Pan-cultural elements in facial displays of emotion. *Science, 164*(3875), 86–88. https://www.science.org/doi/10.1126/science.164.3875.86

[3] Elfenbein, H. A., & Ambady, N. (2002). On the universality and cultural specificity of emotion recognition: A meta-analysis. *Psychological Bulletin, 128*(2), 203–235. https://pubmed.ncbi.nlm.nih.gov/11931516/

[4] Marsh, A. A., Elfenbein, H. A., & Ambady, N. (2003). Nonverbal "accents": Cultural differences in facial expressions of emotion. *Psychological Science, 14*(4), 373–376. https://pubmed.ncbi.nlm.nih.gov/12807413/

[5] Dailey, M. N., et al. (2010). Evidence and a computational explanation of cultural differences in facial expression recognition. *Emotion, 10*(6), 874–893. https://pmc.ncbi.nlm.nih.gov/articles/PMC7360061/

[6] Darwin, C. (1872). *The Expression of the Emotions in Man and Animals.* John Murray. https://www.gutenberg.org/ebooks/1227

[7] Ekman, P. (1972). Universals and cultural differences in facial expressions of emotion. *Nebraska Symposium on Motivation, 19,* 207–283. https://psycnet.apa.org/record/1973-07956-001

[8] Elfenbein, H. A. (2013). Nonverbal dialects and accents in facial expressions of emotion. *Emotion Review, 5*(1), 90–96. https://journals.sagepub.com/doi/abs/10.1177/1754073912451332

[9] Li, R., et al. (2025). Why do cultures affect facial emotion perception? A systematic review. *Journal of Cross-Cultural Psychology.* https://journals.sagepub.com/doi/10.1177/00220221251334811

[10] Tolosana, R., et al. (2020). Deepfakes and beyond: A survey of face manipulation and fake detection. *Information Fusion, 64,* 131–148. https://www.sciencedirect.com/science/article/pii/S1566253520303110

[11] Verdoliva, L. (2020). Media forensics and deepfakes: An overview. *IEEE Journal of Selected Topics in Signal Processing, 14*(5), 910–932. https://ieeexplore.ieee.org/document/9133404

[12] Li, Y., & Lyu, S. (2019). Exposing DeepFake videos by detecting face warping artifacts. *CVPR Workshop.* https://arxiv.org/abs/1811.00656

[13] Chandra, N., et al. (2025). DeepFake-Eval-2024: A large-scale multi-modal benchmark. arXiv preprint. https://www.emergentmind.com/topics/deepfake-eval-2024

[14] Mori, M. (1970/2012). The uncanny valley. *IEEE Robotics & Automation Magazine, 19*(2), 98–100. https://ieeexplore.ieee.org/document/6213181

[15] Elfenbein, H. A., & Ambady, N. (2002). Is there an in-group advantage in emotion recognition? *Psychological Bulletin, 128*(2), 243–249. https://pubmed.ncbi.nlm.nih.gov/11931518/

[16] Elfenbein, H. A., et al. (2007). Toward a dialect theory: Cultural differences in the expression and recognition of posed facial expressions. *Emotion, 7*(1), 131–146. https://psycnet.apa.org/record/2007-02390-010

[17] Elfenbein, H. A., & Ambady, N. (2003). When familiarity breeds accuracy: Cultural exposure and facial emotion recognition. *Journal of Personality and Social Psychology, 85*(2), 276–290. https://psycnet.apa.org/record/2003-06075-007

[18] Matsumoto, D. (1989). Cultural influences on the perception of emotion. *Journal of Cross-Cultural Psychology, 20*(1), 92–105. https://journals.sagepub.com/doi/10.1177/0022022189201006

[19] Matsumoto, D., & Ekman, P. (1989). American-Japanese cultural differences in intensity ratings of facial expressions of emotion. *Motivation and Emotion, 13*(2), 143–157. https://link.springer.com/article/10.1007/BF00992959

[20] Gendron, M., et al. (2014). Perceptions of emotion from facial expressions are

not culturally universal: Evidence from a remote culture. *Emotion, 14*(2), 251–262. https://affective-science.org/pubs/2014/gendron-himba-emotion-2014.pdf

[21] Jack, R. E., et al. (2012). Facial expressions of emotion are not culturally universal. *PNAS, 109*(19), 7241–7244. https://www.pnas.org/doi/10.1073/pnas.1200155109

[22] Yan, X., Young, A. W., & Andrews, T. J. (2017). Cross-cultural similarities and differences in the perception of facial expressions. Springer. https://www-users.york.ac.uk/~ta505/YanYoungAndrews_FaceProcessingReview2017.pdf

[23] Birrer, A., & Just, N. (2025). What we know and don't know about deepfakes. *New Media & Society.* https://journals.sagepub.com/doi/10.1177/14614448241253138

[24] Chesney, R., & Citron, D. K. (2019). Deep fakes: A looming challenge for privacy, democracy, and national security. *California Law Review, 107*(6), 1753–1820. https://www.californialawreview.org/print/deep-fakes-a-looming-challenge-for-privacy-democracy-and-national-security/

[25] Weikmann, T., Greber, H., & Nikolaou, A. (2025). After deception: How falling for a deepfake affects media experience. *Journalism & Mass Communication Quarterly.* https://journals.sagepub.com/doi/10.1177/19401612241233539

[26] Barca, L., et al. (2025). Deepfake media forensics: Status and future challenges. *Applied Sciences, 15*(6). https://pmc.ncbi.nlm.nih.gov/articles/PMC11943306/

[27] Sciancalepore, S., et al. (2025). DeepFake detection in the AIGC era: A survey. *Information Sciences, 719.* https://www.sciencedirect.com/science/article/abs/pii/S1566253525008024

[28] Krumhuber, E. G., et al. (2021). Human and machine recognition of facial expressions across cultures. *Frontiers in Psychology, 12.* https://www.frontiersin.org/articles/10.3389/fpsyg.2021.645404

[29] Zhang, H., et al. (2025). Modeling cultural bias in facial expression recognition with adaptive agents. arXiv preprint. https://arxiv.org/pdf/2510.13557

[30] Hofstede, G. (1984). *Culture's Consequences.* SAGE Publications. https://us.sagepub.com/en-us/nam/cultures-consequences/book6022

[31] MacDorman, K. F., & Kageki, N. (2012). The uncanny valley: The original essay by Masahiro Mori. *IEEE Robotics & Automation Magazine, 19*(2), 98–100. https://ieeexplore.ieee.org/document/6213181

— End of Document —

Inventions Patents Deepfake Detection Ai Avatar Osint Tool

Edit profile

Written by Berend Watchus

65 followers · 27 following

AI & Cyber Sec. Independent Researcher (Non-Profit) exploring AI & interdisciplinary solutions. Former futurist event organizer & care innovation lab volunteer.

No responses yet

Berend Watchus

What are your thoughts?

More from Berend Watchus

In OSINT Team by Berend Watchus

I Built a Real Autonomous AI Researcher (2025)— And Then a Scientist Tried to Rewrite the Timeline...

Author: Berend Watchus. Independent non profit AI & Cyber Security Researcher. March 24, 2026 [Publication for: OSINT TEAM, online...

2d ago 102

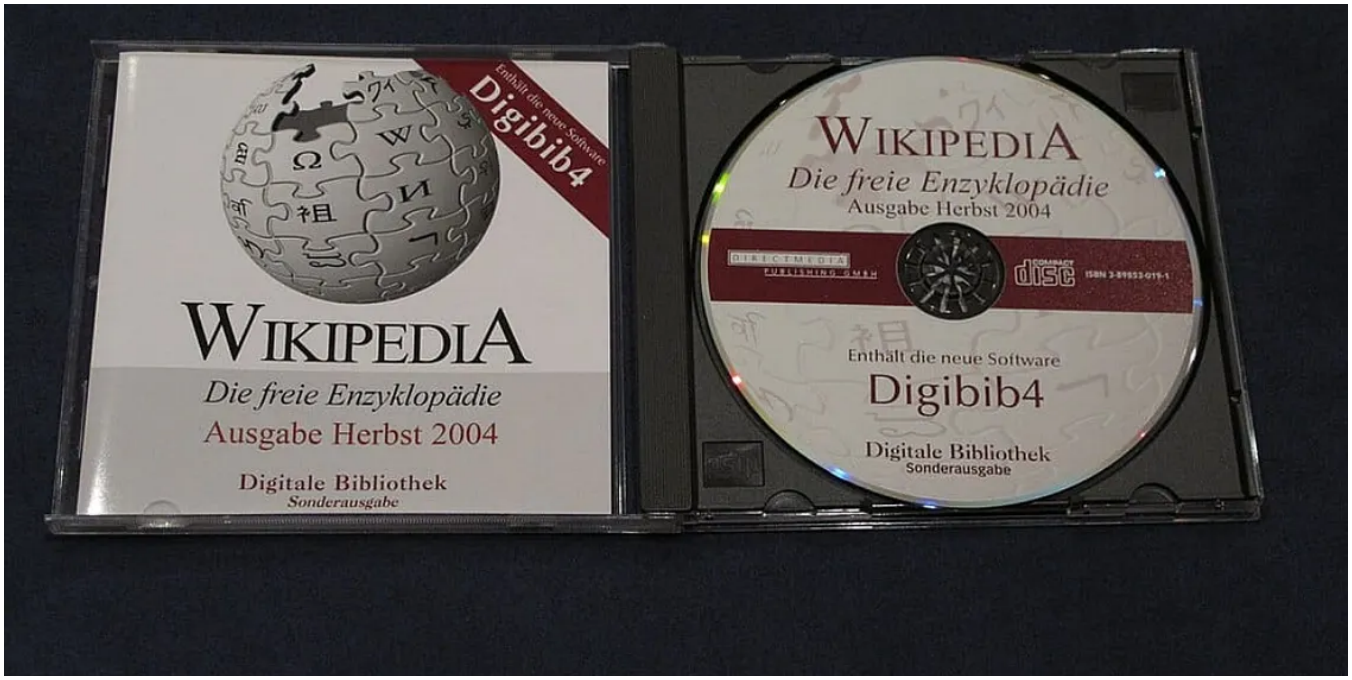

In OSINT Team by Berend Watchus

The Knowledge Base of Humanity Is Compromised — and Nobody Is Minding the Door

Author: Berend Watchus. Independent non profit AI & Cyber Security Researcher. [Publication for: OSINT TEAM, online magazine]

Mar 20 84

In OSINT Team by Berend Watchus

The Outage That Came at Exactly the Right Moment: Claude, Trump, and the PR Value of Good Timing

Author: Berend Watchus. Independent non profit AI & Cyber Sec Researcher. Trendwatcher. [Publication for OSINT TEAM, online magazine]

Mar 16 104

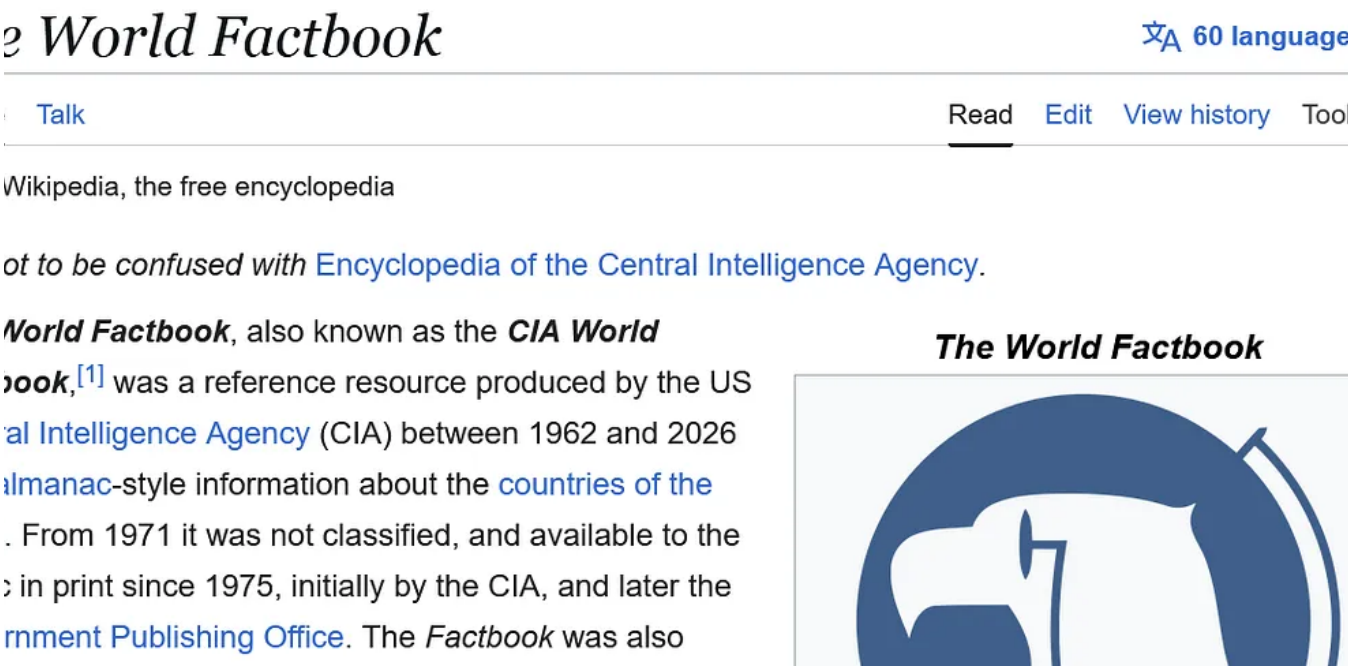
e World Factbook
60 language
Talk
Read
Edit
View history
Tool
Wikipedia, the free encyclopedia
ot to be confused with Encyclopedia of the Central Intelligence Agency.
Vorld Factbook, also known as the CIA World
ook,[1] was a reference resource produced by the US
al Intelligence Agency (CIA) between 1962 and 2026
almanac-style information about the countries of the
. From 1971 it was not classified, and available to the
in print since 1975, initially by the CIA, and later the
rnment Publishing Office. The Factbook was also
The World Factbook

Berend Watchus

THE CIA WORLD FACTBOOK IS GONE — AND YOUR OTHER SOURCES ARE NEXT

Author: Berend F. Watchus — Independent AI & Cybersecurity Researcher (Non-Profit) March 26, 2026 [Publication for: OSINT Team]

1d ago 3

See all from Berend Watchus

Recommended from Medium

In CodeX by MayhemCode

You Are Probably Buying the Wrong Machine for Local AI — The Mac Mini M4 vs Mini PC Truth Nobody...

There is a quiet war happening on the desks of AI enthusiasts right now. On one side, the Mac Mini M4 sits sleek and silent, promising...

6d ago 555 11

In Towards AI by Gao Dalie (高達烈)

NVIDIA Nemoclaw + OpenShell: FASTEST Way to Install

If you don’t have a Medium subscription, use this link to read the full article: link

5d ago 125 1

In Data Science Collective by Florent Poux, Ph.D.

How to Build a Python GUI for 3D Scene Labeling

Learn to label 3D point clouds from 2D image paintings using Python, KD-tree fusion, and DepthAnything v3 in this step-by-step tutorial (+...

281 2

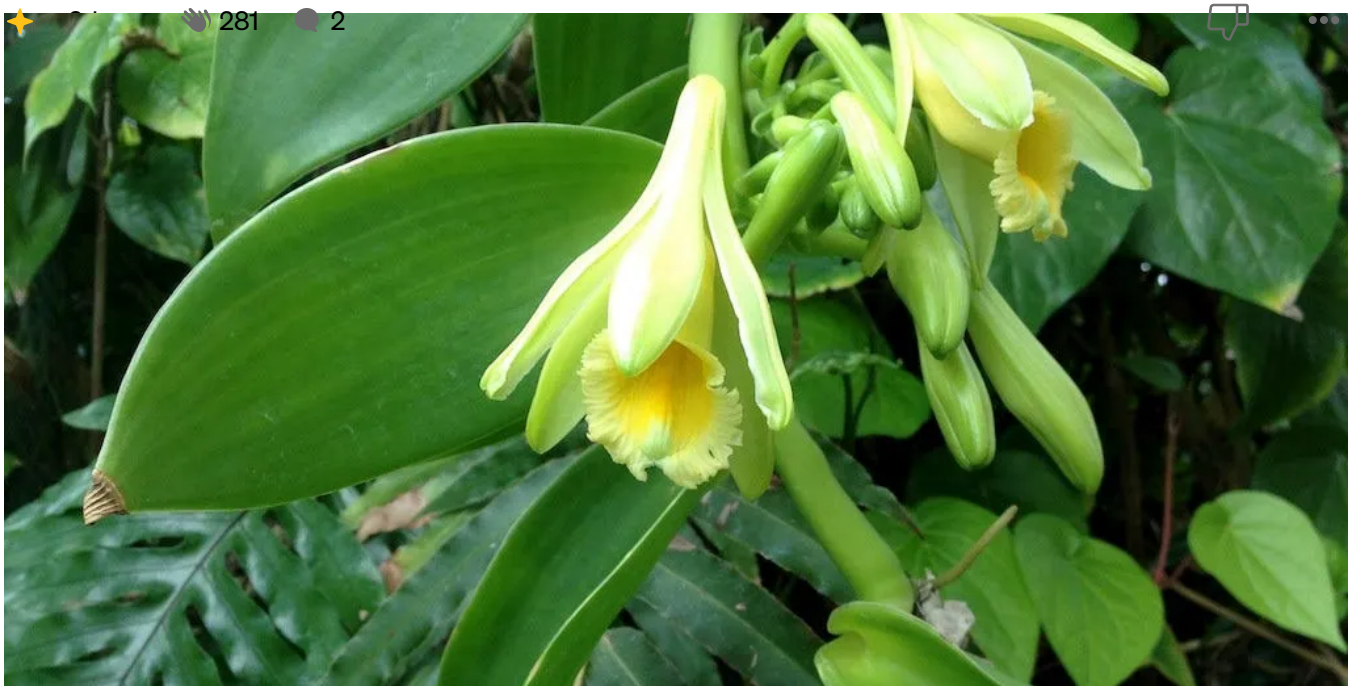

In Science Spectrum by John Knight, PhD

Vanilla Is Anything but Boring

Why a “plain” kitchen staple requires human hands, a lot of work, and a breakout from a chemical vault

3d ago 697 7

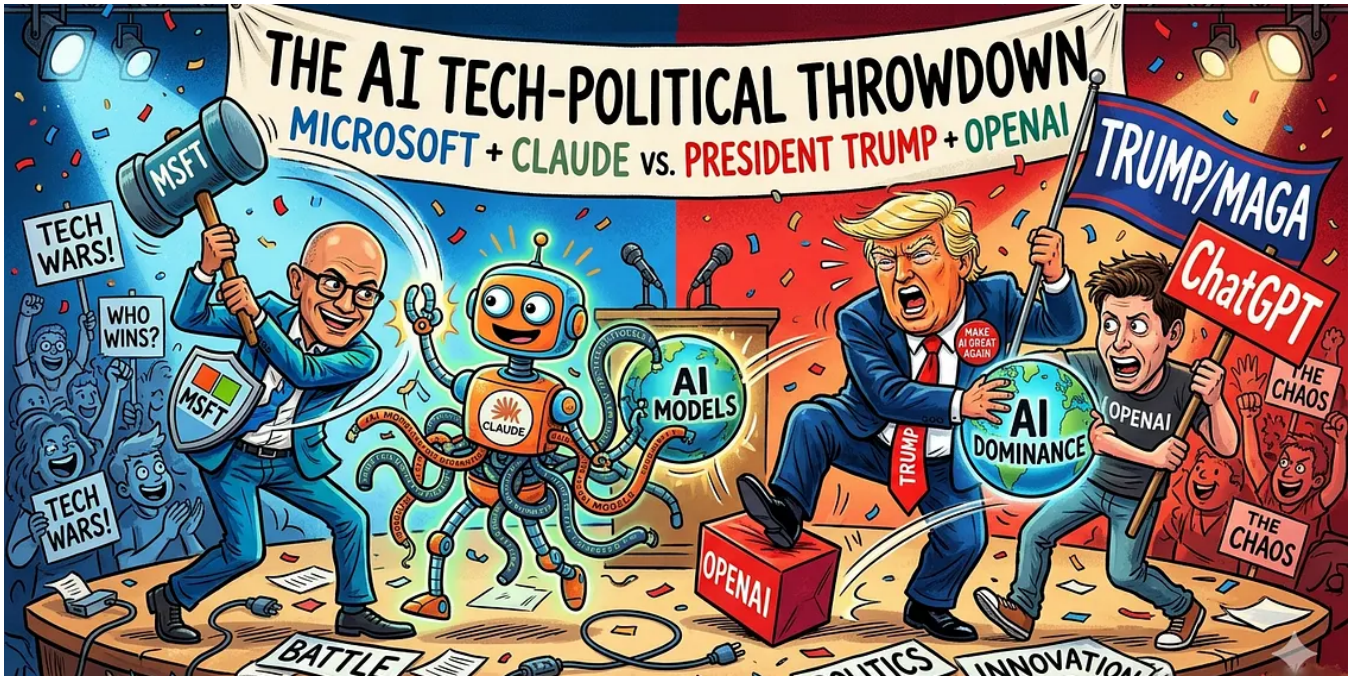

In GitBit by John Gruber

Microsoft Just Picked A Fight With One Of Its Largest Customers, President Trump

When a company known for corporate caution decides to swing at the Pentagon, something deeper is happening beneath the surface.

Mar 16 549 12

In AI Advances by Nov Tech

This AI Improved Itself 100 Times Overnight. Nobody Stopped It.

MiniMax M2.7 didn't just pass the PhD benchmark. It trained for it alone.

4d ago 546 6

See more recommendations